

Large Language Models for Marine Protected Area Design and Management Planning

Ysobel Sims^{1,*}, Arthur Tuda², and Joseph Maina¹

¹School of Natural Sciences, Macquarie University, Macquarie Park, New South Wales, Australia

²Western Indian Ocean Marine Science Association, Zanzibar, Tanzania

*Correspondence: Ysobel Sims (ysobel.sims@mq.edu.au)

KEYWORDS

Marine conservation, marine biodiversity, machine learning, life below water, global biodiversity framework, conservation informatics

Abstract

Achieving biodiversity targets requires advances in marine conservation planning and action. We make three arguments about large language models (LLMs): they add more value when embedded within decision-support systems; they integrate fragmented expertise, data, and tools; they can increase capacity in under-resourced settings through co-design. We illustrate these arguments with a conceptual LLM-enabled marine protected area design and management framework, and address concerns about data, governance, and knowledge equity.

Introduction

Global biodiversity loss and marine ecosystem degradation persist despite decades of international conservation efforts [1, 2, 3]. The Aichi Biodiversity Targets [4] fell short in part because of fragmented implementation, limited capacity, and inadequate monitoring and data systems [5, 6]. The Kunming-Montreal Global Biodiversity Framework (GBF) [7] responds with stronger emphasis on technology, capacity building, and accessible biodiversity knowledge through Targets 20 and 21. Meeting these targets requires advances in conservation planning, monitoring, and decision-making.

Artificial intelligence (AI), including large language models (LLMs), is increasingly being explored for conservation, environmental modelling, and policy [8, 9, 10, 11, 12, 13, 14]. As outlined in Box 1 and Figure 1, LLMs are trained on predictive text at scale, supporting information retrieval, knowledge synthesis, and interaction with analytical tools, with applications ranging from international agreement question answering to coastal risk prioritisation and GIS analysis [15, 16, 17, 18]. However, integrating these capabilities into complete conservation planning workflows remains an open challenge.

Marine protected area (MPA) planning provides an example case. It requires the integration of ecological data, spatial analyses, fisheries science, governance objectives, and policy, often distributed across separate tools, datasets, and institutions. When these inputs are difficult to combine, planners have less capacity to assess ecological trade-offs and adapt management to changing environmental conditions, threats, and policy priorities. These constraints are particularly relevant in under-resourced contexts, where limited expertise and infrastructure can restrict access to integrated evidence and planning tools.

We argue that LLMs can support MPA design and management planning by coordinating established analytical tools, synthesising their outputs, and making evidence accessible to planners and stakeholders. Their role is to support quantitative analysis and human judgement, with planning decisions remaining subject to domain expertise, local context, and appropriate governance. Co-designed with end users, such systems could also broaden participation and support planning in contexts with limited analytical capacity. We develop this argument through the GBF's emphasis on monitoring, planning, and knowledge and capacity building, before presenting a framework for LLM-enabled MPA design and management planning under Target 3. The conceptual framework specifies a workflow from data ingestion to draft management plan, identifies where LLM support is appropriate and where human authority must be retained, and examines the implementation challenges and evaluation needs that must be addressed when developing the system.




Figure 1: Artificial neural networks and machine learning.

(a) Biological neurons transmit signals through synapses, with learning arising from changes in synaptic strength. Artificial neural networks (ANNs) are loosely inspired by this organisation but operate as abstract computational systems rather than biological analogues.

(b) ANNs consist of interconnected nodes whose weighted connections are adjusted during training to improve performance on data-driven tasks.

(c) Large language models (LLMs) are ANNs trained primarily to predict the next token in a sequence from its preceding context. Their large numbers of parameters enable them to capture complex statistical structure in language. Models trained on broad datasets in this way are commonly termed foundation models. Systems such as ChatGPT can also process images and other data types through specialised components, but effective analysis of ecological imagery and remote sensing products generally requires specialised models working alongside the LLM.

(d) Computer vision methods analyse visual information, including photographs and underwater imagery, to identify species and ecological features. Convolutional neural networks are a common approach, learning visual features that support tasks such as species classification.

Informing Biodiversity Conservation under the GBF

The GBF requires conservation decisions to draw on biodiversity evidence, policy objectives, and stakeholder knowledge. The proposed framework uses LLMs to help connect these inputs with the specialised tools used to analyse them. Figure 2 illustrates the framework's architecture, in which the LLM retrieves and synthesises information, translates requests into analytical tasks, and coordinates tools for spatial analysis, modelling, and scenario exploration. Specialised systems perform quantitative analyses, while planners and other authorised stakeholders retain decision authority. Three areas of GBF implementation inform the framework: analysis and monitoring, planning and decision-making, and knowledge and capacity building.




Figure 2: The potential of LLMs is apparent when considering their central role in a wider system. Tools such as computer vision models, geospatial models and forecasting models are used by the LLM to provide conservation decision support guided by stakeholder interaction. Stakeholders ensure appropriate governance, supporting equitable outcomes that progress towards biodiversity targets.

Analysis and Monitoring

Analysis and monitoring provide evidence for MPA design and management, including ecosystem assessment, spatial planning, fisheries management, and restoration [5]. LLM-enabled systems can help connect observations from visual, acoustic, and spatial data with specialised analytical models. Multimodal and task-specific models can support ecological classification [19, 20, 21], while robotic and autonomous systems can extend data collection in difficult-to-access environments [22, 23]. These systems provide complementary functions: LLMs can help retrieve, organise, and route evidence through analytical workflows, while specialised models remain responsible for tasks such as ecological classification and quantitative analysis.

LLMs can also translate natural-language requests into geospatial workflows and coordinate pipelines for spatial modelling, forecasting, and data harmonisation [24, 25, 26]. This can make analytical workflows more accessible and connect their outputs to subsequent planning tasks, without replacing the specialist methods used to generate the evidence.

Planning and Decision-Making

The GBF calls for biodiversity to be incorporated into decision-making (Target 14) and for sustainable management of fisheries and aquaculture (Target 10). Implementing these targets requires ecological evidence, policy constraints, and management priorities to be considered together across spatial and temporal scales.

LLMs can support this process by retrieving relevant information, synthesising evidence, and helping planners explore alternative scenarios. Existing applications include question answering over international agreements [15], scenario evaluation [16], and co-ordination of agents with competing land-use priorities within a spatial model [17]. In MPA planning, these functions can help translate policy objectives and management constraints into analytical tasks, organise alternative configurations, and communicate their implications and trade-offs. Quantitative evaluation depends on the underlying models and decision criteria, while the assessment of trade-offs and selection of management actions remain with planners and other authorised decision-makers.

Knowledge and Capacity Building

MPA design and management planning draws on expertise across ecology, fisheries science, spatial analysis, governance, and policy. Relevant information is distributed across scientific publications, policy documents, multilingual sources, and grey literature [8, 27]. Within the proposed framework, domain-specific retrieval can improve access to this material [28, 29], while natural-language interfaces can help planners and stakeholders engage with GIS and modelling tools without requiring equivalent proficiency in each technical domain [18, 30]. By connecting information and analytical capabilities within the planning workflow, the framework could support participation by stakeholders with different forms of expertise, including in under-resourced contexts where specialist capacity and infrastructure are limited.

The framework must also account for knowledge governance. Under GBF Targets 20–22, capacity building and access to biodiversity knowledge must respect the rights and authority of Indigenous Peoples and local communities. The inclusion of Indigenous and locally held knowledge in MPA planning requires the consent of its custodians and governance arrangements that reflect their authority over its use.




Figure 3: The proposed architecture for an LLM-assisted MPA design and management planning tool. Inputs are harmonised by the LLM system and used in spatial prioritisation, scenario design and plan generation modules. The tool outputs rationale, site and zoning proposals, and draft documents. The user can explore the data, make selections, and use natural language to modify outputs.

LLM-Enabled MPA Design and Management Planning

The capabilities described above come together in MPA design and management planning, supporting site selection, zoning, and management plan development under GBF Target 3. Achieving the target requires protected areas that are ecologically representative and connected, while accounting for social, economic, and governance considerations. LLMs can provide a coordination layer that links these inputs and analytical outputs within an iterative planning workflow.

MPA planning integrates ecological data, spatial analyses, governance objectives, stakeholder priorities, and legal frameworks [5], often distributed across separate tools and institutions [31]. Figure 3 presents a conceptual architecture for connecting these components through knowledge synthesis, scenario exploration, spatial prioritisation, and plan generation. The LLM translates between users and specialised tools, prepares analytical inputs, interprets outputs, and communicates alternatives, while users retain control over consequential decisions. Box 2 specifies the proposed workflow.

Box 2: Proposed workflow for LLM-assisted MPA design and management planning.

This box describes the intended workflow of the proposed framework. At each step, the LLM exchanges structured data with specialised components, which perform the retrieval, spatial, and optimisation tasks.

Step 1: Data ingestion and requirements interpretation. The user uploads local data (species observations, habitat maps, community boundaries, use zones). The system retrieves national and international layers (remote sensing products, GBIF records) through data services, and a retrieval-augmented generation (RAG) component retrieves passages from national biodiversity strategies, legislation, and literature. Each dataset and passage is tagged at ingestion with provenance metadata recording its source, custodian, and conditions of use. The LLM identifies applicable GBF targets, national obligations, and planning constraints. Culturally sensitive layers remain under the control of their custodians and enter the analysis only in forms they approve. The user confirms data and requirements before analysis.

Step 2: Spatial prioritisation. GIS tools harmonise layers to a common set of planning units. The LLM translates requirements into a structured problem specification (conservation features, planning units, targets, costs, and constraints) and mathematical formulation (e.g., Mixed Integer Linear Programming or Maximum Coverage Problem) for an established software suit such as prioritizr or Marxan, which performs the optimisation. The solver returns candidate configurations, and the LLM generates a structured rationale identifying feature representation, trade-offs, and key assumptions.

Step 3: User review of candidate site configurations. The user explores candidate maps and rationale, queries results, and adjusts inputs (targets, exclusions, community knowledge). The LLM updates and compares prioritisation scenarios. The user confirms site boundaries.

Step 4: Zonation and scenario analysis. The workflow is applied to internal zoning, incorporating use-type data (fishing, tourism, cultural sites) and legal constraints. Scenario analysis evaluates alternative configurations under different assumptions, with trade-offs summarised in accessible language.

Step 5: User review of candidate zone configurations. The user reviews zoning options, adjusts inputs as needed, and selects a preferred configuration.

Step 6: Management plan generation. The LLM drafts a management plan from the confirmed boundary, zoning, stakeholder inputs, and budget, following national guidelines and precedent documents retrieved through the RAG component. The user reviews and refines the draft before export.

Throughout all stages, provenance metadata is carried through each analytical step to every output, so each map, rationale, and draft plan identifies its data sources, custodians, and conditions of use. Outputs are also logged with model versions and solver parameters, and the user retains authority over consequential decisions.

Spatial prioritisation uses optimisation methods such as those implemented in Marxan and prioritizr. The LLM translates stakeholder objectives into analytical tasks, prepares inputs, interprets outputs, and communicates alternatives. The framework could also support alternative optimisation algorithms, including NSGA-III, while keeping quantitative optimisation separate from the LLM's role in natural language interaction. This allows planners to explore multiple configurations and understand their trade-offs without requiring the LLM to perform the optimisation itself.

The framework also supports adaptive management by connecting monitoring and analytical systems to the planning workflow. Species observations, geospatial updates, forecasting outputs, and robotic sensing can trigger a review of existing plans, creating a feedback loop between monitoring, analysis, and planning. This allows plans to be

updated as ecological conditions, threats, and management priorities change.

Uncertainty arises from incomplete data, forecasting error, model variation, and conflicting priorities. Scenario analysis and provenance tracking can help address these sources of uncertainty. Solver runs across alternative targets, costs, and input datasets can identify the stability of candidate configurations, while source-linked rationales help distinguish evidence-supported statements from unsupported outputs. Results can therefore report the sensitivity of planning outcomes and identify where additional monitoring would improve future decisions.

Protected area design also involves normative judgements about ecological priorities, resource use, and acceptable trade-offs. LLMs can structure information, explore alternatives, coordinate analytical systems, and communicate rationale, but conservation objectives and decisions remain human responsibilities. User review and decision authority are therefore explicit components of the workflow, rather than post hoc safeguards. The intended outcome is to increase planning capacity and support more consistent integration of ecological, socioeconomic, and Indigenous knowledge while retaining accountable human governance.

Evaluating the Framework

The framework is conceptual and requires empirical validation before deployment. Evaluation should operate at four levels. At the component level, retrieval should be assessed for accuracy and source attribution against expert-curated question sets, while LLM-generated problem specifications should be compared with those produced by experienced planners using the same objectives and data. At the system level, repeated runs and model updates should be tested for reproducibility, with uncertainty tracked from inputs to user-facing outputs. At the user level, co-design studies with practitioners and communities should assess usability, interpretation of rationale and uncertainty, and whether decision authority remains with users in practice. At the outcome level, plans developed with and without the system should be compared against ecological criteria such as habitat representation and connectivity and, where feasible, indicators of management effectiveness.

Once validated, the system could also support the evaluation of individual plans. Before implementation, conservation planning guidelines, IUCN standards, and GBF targets can provide structured criteria for identifying gaps in draft plans. After implementation, monitoring data can feed back into the system to assess progress towards objectives and identify where revision is needed. LLM-guided assessment could apply these criteria at scale, but would itself require validation against expert review. At every level, expert judgement remains the reference standard for assessing system outputs.

Implementation Challenges

Integrating heterogeneous data, analytical models, and human decisions raises challenges for reliability, governance, and knowledge representation. These challenges are important because MPA planning affects conservation priorities, resource allocation, and stakeholder and Indigenous rights.

Data Limitations and Reliability

System performance depends on the quality and coverage of environmental data and the behaviour of individual model components. Off-the-shelf LLMs rarely achieve best-in-class

performance in domain-specific applications without adaptation [19, 20, 21]. They should therefore coordinate specialised ecological, spatial, and optimisation models rather than replace them. Marine conservation datasets are often sparse, heterogeneous, and unevenly distributed, particularly where biodiversity is high but monitoring capacity is limited [9]. These gaps affect spatial prioritisation and forecasting, so outputs should be interpreted in the context of available data and model assumptions. Sensitivity analyses can identify how candidate plans change under alternative datasets, targets, and costs.

LLM variability also affects reproducibility. External models may be updated or replaced, and outputs can vary between runs. Version-controlled deployment or local hosting can reduce this variability, with model updates evaluated before being incorporated into decision-support workflows. Logging model versions, solver parameters, inputs, and outputs, as specified in Box 2, allows results to be traced to the system configuration that produced them.

Governance, Transparency and Accountability

Governance begins with data access, ownership, and consent. This is particularly important for Indigenous and locally held knowledge, where misuse can affect participation and legitimacy. The CARE principles for Indigenous data governance (collective benefit, authority to control, responsibility, and ethics) [32] provide a basis for determining how such knowledge enters the system. Provenance metadata can record custodianship and conditions of use, while custodians determine whether sensitive data enter the analysis and in what form. For example, locations may be generalised or excluded from externally hosted models.

Human oversight should be built into the workflow rather than left to a final review step. Practitioners should be able to interrogate assumptions, refine inputs, and adjust outputs [33, 34]. Stakeholder review has corrected model-generated risk assessments in environmental decision-support applications [16], while supervised workflows can outperform fully automated alternatives [14]. The workflow therefore reserves consequential decisions, including confirmation of requirements, site boundaries, zoning, and management plans, for authorised users. Outputs should expose their assumptions, limitations, and supporting sources, with decision logs recording who selected each option and on what basis. MPA stages involving institutional legitimacy, such as mandate setting, stakeholder negotiation, zoning decisions, and formal approval, require human authority [35, 36].

Representation and Knowledge Equity

LLMs reflect biases and gaps in their training data, including uneven representation of ecological knowledge systems and Western-centric assumptions [37, 12, 38]. These biases can influence how threats are framed, which ecosystems are prioritised, and how local and Indigenous knowledge is interpreted. Bias can enter through data selection, model development, interface design, and deployment [39], potentially affecting prioritisation outcomes under GBF Targets 20–23.

Co-design with practitioners and communities can help identify inappropriate assumptions and align the system with existing planning processes. Language accessibility and multilingual interfaces affect whose knowledge can enter the planning process and who can participate. The system should also allow users to compare outputs under different datasets, weighting schemes, and stakeholder-defined objectives, rather than presenting a single result as neutral. These measures do not eliminate representational bias, but they

make its effects more visible and give stakeholders greater control over how knowledge shapes planning outcomes.

Conclusion

Marine conservation planning is constrained by fragmented data, expertise, and analytical tools, as well as the need to integrate ecological, socioeconomic, and governance considerations under uncertainty. When embedded within conservation workflows and integrated with established analytical tools, LLMs can help address this fragmentation by connecting heterogeneous information, coordinating analytical workflows, and making outputs accessible to planners and stakeholders. The proposed framework outlines how these functions could support MPA planning from data integration to management plan development, while expanding access to conservation knowledge and retaining human decision authority.

The framework remains conceptual and requires empirical evaluation. Co-design with practitioners and communities across different institutional and regional contexts can establish whether these systems improve planning processes and support more ecologically representative and durable plans. Evaluation should also examine reliability, reproducibility, knowledge governance, and the effects of model and data limitations. Evidence across these dimensions is needed to establish where LLM-enabled MPA planning can contribute to biodiversity conservation.

Acknowledgements

No funding was received for this research.

Author Contributions

YS led the manuscript and contributed expertise in machine learning and large language models. AT contributed expertise in applied contexts and marine protected areas, with particular emphasis on the presented use cases. JM contributed expertise in marine conservation. All authors contributed to the conceptualisation and writing of the manuscript.

Competing Interests

The authors declare no competing financial or non-financial interests.

References

- [1] United Nations Environment Programme. Adaptation gap report 2024: Come hell and high water. Technical report, United Nations Environment Programme, Nairobi, Kenya, November 2024.
- [2] United Nations Environment Programme. Emissions gap report 2024: Narrowing the gap. Technical report, United Nations Environment Programme, Nairobi, Kenya, November 2024.

- [3] World Wide Fund for Nature. Living planet report 2024 – a system in peril. Technical report, WWF, Gland, Switzerland, October 2024.
- [4] Secretariat of the Convention on Biological Diversity. Strategic plan for biodiversity 2011–2020 and the Aichi biodiversity targets. Technical report, Convention on Biological Diversity, Montreal, Canada, 2010.
- [5] Secretariat of the Convention on Biological Diversity. Global biodiversity outlook 5. Technical report, Convention on Biological Diversity, Montreal, Canada, 2020.
- [6] Haigen Xu, Yiming Cao, Dongping Yu, Min Cao, Yafei He, Margareth Gill, and Henrique M. Pereira. Ensuring effective implementation of the post-2020 global biodiversity targets. *Nature Ecology & Evolution*, 5(4):411–418, 2021.
- [7] Secretariat of the Convention on Biological Diversity. Kunming-Montreal global biodiversity framework. Technical report, Convention on Biological Diversity, Montreal, Canada, 2022.
- [8] Sam A. Reynolds, Sara Beery, Neil Burgess, Mark Burgman, Stuart H.M. Butchart, Steven J. Cooke, David Coomes, Finn Danielsen, Enrico Di Minin, América Paz Durán, Francis Gassert, Amy Hinsley, Sadiq Jaffer, Julia P.G. Jones, Binbin V. Li, Oisín Mac Aodha, Anil Madhavapeddy, Stephanie A.L. O'Donnell, William M. Oxbury, Lloyd Peck, Nathalie Pettorelli, Jon Paul Rodríguez, Emily Shuckburgh, Bernardo Strassburg, Hiromi Yamashita, Zhongqi Miao, and William J. Sutherland. The potential for AI to revolutionize conservation: a horizon scan. *Trends in Ecology & Evolution*, 40(2):191–207, 2025.
- [9] Laura J. Pollock, Justin Kitzes, Sara Beery, Kaitlyn M. Gaynor, Marta A. Jarzyna, Oisín Mac Aodha, Bernd Meyer, David Rolnick, Graham W. Taylor, Devis Tuia, and Tanya Berger-Wolf. Harnessing artificial intelligence to fill global shortfalls in biodiversity knowledge. *Nature Reviews Biodiversity*, 1(3):166–182, March 2025.
- [10] Christos Mammides, Hao Gu, Thilina S. Nimalrathna, Naufal Rahman Avicena, Harris Papadopoulos, and Ahimsa Campos-Arceiz. Emerging applications of large language models in ecology and conservation science. *Conservation Biology*, 2026.
- [11] Chris Tyler, K L Akerlof, Alessandro Allegra, Zachary Arnold, Henriette Canino, Marius A Doornenbal, Josh A Goldstein, David Budtz Pedersen, and William J Sutherland. AI tools as science policy advisers? The potential and the pitfalls. *Nature*, 622(7981):27–30, October 2023.
- [12] Matt Ziegler, Sarah Lothian, Brian O'Neill, Richard Anderson, and Yoshitaka Ota. AI language models could both help and harm equity in marine policymaking. *npj Ocean Sustainability*, 4(1):32, June 2025.
- [13] Nicole DeSantis, Christina Supples, Lea Phillips, Julien Pigot, Jamison Ervin, and Toby Wade. Leveraging AI for enhanced alignment of national biodiversity targets with the global biodiversity goals. *Nature-Based Solutions*, 7:100198, 2025.
- [14] Qiyang Nie and Tong Liu. Large language models for environmental modeling: Framework, capabilities, constraints. *Journal of Environmental Management*, 398: 128417, 2026.

- [15] Xiaowei Wang, Mingdan Zhang, Hao Liu, Xiaodong Ma, Yingchao Liu, and Yitong Chen. ChatBBNJ: a question–answering system for acquiring knowledge on biodiversity beyond national jurisdiction. *Frontiers in Marine Science*, 11, 2024.
- [16] Maria Katherina Dal Barco, Veronica Casartelli, Marcello Sanò, Sebastiano Vascon, Silvia Torresan, and Andrea Critto. Prioritise risks and improve adaptation strategies in the Veneto coast through the application of a custom AI tool. *International Journal of Disaster Risk Reduction*, 130:105818, 2025.
- [17] Y. Zeng, C. Brown, M. Byari, J. Raymond, T. Schmitt, and M. Rounsevell. InsNet-CRAFTY v1.0: integrating institutional network dynamics powered by large language models with land use change simulation. *Geoscientific Model Development*, 18(15): 4983–5013, 2025.
- [18] T. Akinboyewa, Z. Li, H. Ning, and M. N. Lessani. GIS Copilot: towards an autonomous GIS agent for spatial analysis. *International Journal of Digital Earth*, 18(1), 2025.
- [19] Yijia Shun Qi, Shuzhang Cai, Zunduo Zhao, Jiaming Li, Yanbin Lin, and Zhiqiang Wang. Benchmarking large language models for image classification of marine mammals. In *2024 IEEE International Conference on Knowledge Graph (ICKG)*, pages 258–265, 2024.
- [20] Jiaxin Zhu, Shibai Yin, Xin Liu, Xingyang Wang, and Yee-Hong Yang. FishDetectLLM: Multimodal instruction tuning with large language models for fish detection. *Knowledge-Based Systems*, 318:113418, 2025.
- [21] Hongyong Han, Wei Wang, Gaowei Zhang, Mingjie Li, and Yi Wang. Enhancing vision-language models with morphological and taxonomic knowledge: Towards coral recognition for ocean health. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(27):28052–28060, April 2025.
- [22] Dinesh Jackson Samuel, Yusuf Sermet, David Cwiertny, and Ibrahim Demir. Integrating vision-based AI and large language models for real-time water pollution surveillance. *Water Environment Research*, 96(8), 2024.
- [23] Waseem Akram, Muhayy Ud Din, Abdelhaleem Saad, and Irfan Hussain. AquaChat: An LLM-guided ROV framework for adaptive inspection of aquaculture net pens. *Aquacultural Engineering*, 111:102607, 2025.
- [24] Hartwig H. Hochmair, Levente Juhász, and Takoda Kemp. Correctness comparison of ChatGPT-4, Gemini, Claude-3, and Copilot for spatial tasks. *Transactions in GIS*, 28(7):2219–2231, 2024.
- [25] Caichao Lv, Ning Song, Jie Nie, Min Ye, Xinyue Liang, Xiong Liu, and Jun Ma. MS-LIP: Multiscale semantic information integration with large language models for marine prediction. *IEEE Transactions on Geoscience and Remote Sensing*, 63:1–15, 2025.
- [26] Shuyu Wang, Wengen Li, Hanchen Yang, Jihong Guan, Xiwei Liu, Yichao Zhang, Rufu Qin, and Shuigeng Zhou. LLM4HRS: LLM-based spatiotemporal imputation model for highly sparse remote sensing data. *IEEE Transactions on Geoscience and Remote Sensing*, 63:1–17, 2025.

- [27] Oded Berger-Tal, Bob B.M. Wong, Carrie Ann Adams, Daniel T. Blumstein, Ulrika Candolin, Matthew J. Gibson, Alison L. Greggor, Malgorzata Lagisz, Biljana Macura, Catherine J. Price, Breanna J. Putman, Lysanne Snijders, and Shinichi Nakagawa. Leveraging AI to improve evidence synthesis in conservation. *Trends in Ecology & Evolution*, 39(6):548–557, 2024.
- [28] Nabin Khanal, Chun Meng Yu, Jui-Cheng Chiu, Anav Chaudhary, Ziyue Zhang, Kakani Katija, and Angus G. Forbes. FathomGPT: A natural language interface for interactively exploring ocean science data. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*, pages 1–15, New York, NY, USA, 2024. Association for Computing Machinery.
- [29] Elise C. Gallois, Arianna Salili-James, Sanson T. S. Poon, Artur Trebski, and David W. Redding. Fast-tracking ecological interpretation using bespoke quantitative large language models. *Methods in Ecology and Evolution*, 16(12):2730–2740, 2025.
- [30] Hartwig H. Hochmair. Use and effectiveness of chatbots as support tools in GIS programming course assignments. *ISPRS International Journal of Geo-Information*, 14(4), 2025.
- [31] Kirsten Grorud-Colvert, Jenna Sullivan-Stack, Callum Roberts, Vanessa Constant, Barbara Horta e Costa, Elizabeth P. Pike, Naomi Kingston, Dan Laffoley, Enric Sala, Joachim Claudet, Alan M. Friedlander, David A. Gill, Sarah E. Lester, Jon C. Day, Emanuel J. Gonçalves, Gabby N. Ahmadi, Matt Rand, Angelo Villagomez, Natalie C. Ban, Georgina G. Gurney, Ana K. Spalding, Nathan J. Bennett, Johnny Briggs, Lance E. Morgan, Russell Moffitt, Marine Deguignet, Ellen K. Pikitch, Emily S. Darling, Sabine Jessen, Sarah O. Hameed, Giuseppe Di Carlo, Paolo Guidetti, Jean M. Harris, Jorge Torre, Zafer Kizilkaya, Tundi Agardy, Philippe Cury, Nirmal J. Shah, Karen Sack, Ling Cao, Miriam Fernandez, and Jane Lubchenco. The MPA Guide: A framework to achieve global goals for the ocean. *Science*, 373(6560), 2021.
- [32] Stephanie Russo Carroll, Ibrahim Garba, Oscar L. Figueroa-Rodríguez, Jarita Holbrook, Raymond Lovett, Simeon Materechera, Mark Parsons, Kay Raseroka, Desi Rodriguez-Lonebear, Robyn Rowe, Rodrigo Sara, Jennifer D. Walker, Jane Anderson, and Maui Hudson. The care principles for indigenous data governance. *Data Science Journal*, November 2020.
- [33] Yunfan Zhao, Niclas Boehmer, Aparna Taneja, and Milind Tambe. Towards foundation-model-based multiagent system to accelerate AI for social impact. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems*, page 2901–2907, Richland, SC, 2025. International Foundation for Autonomous Agents and Multiagent Systems.
- [34] Mark Steyvers, Heliodoro Tejeda, Aakriti Kumar, Catarina Belem, Sheer Karny, Xinyue Hu, Lukas W. Mayer, and Padhraic Smyth. What large language models know and what people think they know. *Nature Machine Intelligence*, 7(2):221–231, February 2025.
- [35] Lee Thomas and Julie Middleton. *Guidelines for Management Planning of Protected Areas*. Number 10 in Best Practice Protected Area Guidelines Series. IUCN, 2003.
- [36] Food and Agriculture Organization of the United Nations. *Guidelines for Protected Area Management Planning*. FAO, Rome, Italy, 2022.

- [37] Danilo Urzedo, Zarrin Tasnim Sworna, Andrew J. Hoskins, and Cathy J. Robinson. Al chatbots contribute to global conservation injustices. *Humanities and Social Sciences Communications*, 11(1):204, February 2024.
- [38] Andrea Santangeli, Stefano Mammola, Veronica Nanni, and Sergio A. Lambertucci. Large language models debunk fake and sensational wildlife news. *Integrative Conservation*, 3(2):127–133, 2024.
- [39] Syed Masiur Rahman, Asif Raihan, and Shadi Abudalfa. Generative artificial intelligence for environmental assessment: A new paradigm for sustainability analysis. *Environmental Management*, 76(3), 2026.